



(12) **United States Patent**  
**Wells et al.**

(10) **Patent No.:** **US 9,360,909 B2**  
(45) **Date of Patent:** **Jun. 7, 2016**

(54) **SYSTEM, METHOD AND APPARATUS FOR ENERGY EFFICIENCY AND ENERGY CONSERVATION BY CONFIGURING POWER MANAGEMENT PARAMETERS DURING RUN TIME**

(2013.01); **G06F 1/3296** (2013.01); **Y02B 60/1217** (2013.01); **Y02B 60/1285** (2013.01)

(58) **Field of Classification Search**

CPC ..... **G06F 1/206**; **G06F 1/24**; **G06F 1/26**; **G06F 1/3203**; **G06F 1/3296**; **G06F 1/324**; **G06F 1/3206**; **Y02B 60/1217**; **Y02B 60/1285**  
See application file for complete search history.

(71) Applicants: **Ryan D. Wells**, Folsom, CA (US); **Sanjeev Jahagirdar**, Folsom, CA (US); **Inder Sodhi**, Folsom, CA (US); **Jeremy Shrall**, Portland, OR (US); **Stephen Gunther**, Beaverton, OR (US); **Daniel Ragland**, Hillsboro, OR (US); **Nicholas Adams**, Beaverton, OR (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,634,131 A 5/1997 Matter et al.  
2005/0192788 A1 9/2005 Hobson  
2009/0228222 A1 9/2009 Fantoni

(Continued)

FOREIGN PATENT DOCUMENTS

CN 1877492 A 12/2006  
CN 101923383 A 12/2010

(Continued)

OTHER PUBLICATIONS

Dobrowolski, Daniel, "Intel DX58SO2 LGA 1366 X58 Motherboard Review", Mar. 28, 2011, HARDOCP, pp. 1-26.\*

(Continued)

(72) Inventors: **Ryan D. Wells**, Folsom, CA (US); **Sanjeev Jahagirdar**, Folsom, CA (US); **Inder Sodhi**, Folsom, CA (US); **Jeremy Shrall**, Portland, OR (US); **Stephen Gunther**, Beaverton, OR (US); **Daniel Ragland**, Hillsboro, OR (US); **Nicholas Adams**, Beaverton, OR (US)

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 437 days.

(21) Appl. No.: **13/843,405**

(22) Filed: **Mar. 15, 2013**

(65) **Prior Publication Data**

US 2013/0283026 A1 Oct. 24, 2013

**Related U.S. Application Data**

(60) Provisional application No. 61/635,817, filed on Apr. 19, 2012.

(51) **Int. Cl.**  
**G06F 1/26** (2006.01)  
**G06F 1/24** (2006.01)  
**G06F 1/32** (2006.01)

(52) **U.S. Cl.**  
CPC .. **G06F 1/26** (2013.01); **G06F 1/24** (2013.01); **G06F 1/3203** (2013.01); **G06F 1/324**


*Primary Examiner* — Dennis M Butler

(74) *Attorney, Agent, or Firm* — Nicholson De Vos Webster & Elliott, LLP

(57) **ABSTRACT**

According to one embodiment of the invention, an integrated circuit device at least one compute engine and a control unit. Coupled to the compute engine(s), the control unit is adapted to dynamically control an energy-efficient operating setting of at least one power management parameter for the integrated circuit device after execution of Basic Input/Output System (BIOS) has already completed.

**24 Claims, 15 Drawing Sheets**



(56)

**References Cited**

TW

201133228 A 10/2011

## U.S. PATENT DOCUMENTS

2009/0235108 A1\* 9/2009 Gold ..... G06F 1/206  
713/500  
2009/0256546 A1\* 10/2009 Wu ..... G06F 1/08  
323/318  
2010/0115304 A1 5/2010 Finkelstein et al.  
2010/0117579 A1\* 5/2010 Culbert et al. .... 318/471  
2010/0122008 A1\* 5/2010 Goss et al. .... 710/264  
2010/0153761 A1 6/2010 Nishioka  
2010/0223489 A1\* 9/2010 Huang ..... G06F 1/08  
713/501  
2010/0281284 A1 11/2010 Kasprzak et al.  
2011/0047396 A1 2/2011 Kawamoto et al.  
2011/0153924 A1 6/2011 Vash et al.  
2011/0161706 A1\* 6/2011 Huang ..... G06F 1/08  
713/323  
2011/0241628 A1 10/2011 Chen  
2011/0258477 A1\* 10/2011 Baker ..... G06F 1/3203  
713/600  
2012/0110352 A1\* 5/2012 Branover ..... G06F 1/206  
713/300  
2012/0216058 A1\* 8/2012 Wells et al. .... 713/322  
2013/0061064 A1 3/2013 Ananthakrishnan et al.

## FOREIGN PATENT DOCUMENTS

EP 0676686 A2 10/1995  
KR 10-2009-0084255 A 8/2009  
TW M412423 U 9/2011

## OTHER PUBLICATIONS

Mohr, Neil; Best Overclocking software: 15 Top Apps; www.techradar.com; Oct. 6, 2010 TechRadar; pp. 1-26.\*  
Office action from U.S. Appl. No. 13/335,635, mailed Jun. 17, 2014, 16 pages.  
Final Office action from U.S. Appl. No. 13/335,635, mailed Jan. 13, 2015, 18 pages.  
Office action with English translation from Taiwan Patent Application No. 101147198, mailed Aug. 7, 2014, 16 pages.  
Office action with English translation from Taiwan Patent Application No. 101147198, mailed Apr. 22, 2015, 24 pages.  
PCT/US2012/068270 Notification of Transmittal of the International Search Report and the Written Opinion of the International Searching Authority, mailed Apr. 1, 2013, 10 pages.  
PCT/US2012/068270 Notification Concerning Transmittal of International Preliminary Report on Patentability, mailed Jul. 3, 2014, 7 pages.  
Office action from U.S. Appl. No. 13/335,635, mailed Aug. 14, 2015, 18 pages.  
Final Office action from U.S. Appl. No. 13/335,635, mailed Jan. 5, 2016, 17 pages.  
Office action with English translation from Chinese Patent Application No. 201280063842.0, mailed Nov. 18, 2015, 23 pages.

\* cited by examiner




FIG. 1




FIG. 2




FIG. 3




FIG. 4A




FIG. 5

FIG. 4B

MSR Address:601h			
Access:Read/Write			
Bit	Scope	Default	Description
63:62	-	-	Reserved
61:59	Package	011b	<b>PSI3_CODE (R/W)</b> PSI Code associated with current threshold specified in PSI3_THRESHOLD field.
58:52	Package	01h	<b>PSI3_THRESHOLD (R/W)</b> Current Threshold for PSI3_CODE specified in Amps
51:49	Package	010b	<b>PSI2_CODE (R/W)</b> PSI Code associated with current threshold specified in PSI2_THRESHOLD field.
48:42	Package	05h	<b>PSI2_THRESHOLD (R/W)</b> Current Threshold for PSI2_CODE specified in Amps
41:39	Package	001b	<b>PSI1_CODE (R/W)</b> PSI Code associated with current threshold specified in PSI1_CURRENT field.
38:32	Package	14h	<b>PSI1_CODE (R/W)</b> Current threshold for PSI1_CODE specified in Amps
31	Package	0b	<b>CURRENT_LIMIT_LOCK (RW/L)</b> When set, locks bits 31:0 of this MSR from further writes until a reset.
30:13	-	-	Reserved
12:0	Package	0x380	<b>CURRENT_LIMIT (RW/L)</b> Maximum instantaneous current allowed at any given time. The value is represented in 1/8 A increments. Current Limit=[Value] * (1/8) <b>Note:</b> System BIOS must program this field with the maximum current supported by the Voltage Regulator.

PRIMARY_PLANE_LL_R_CONFIG_CONTROL <u>452</u>						
MSR Symbol:		MSR Address: 0x603				
Symbol	Name	MSB:LSB	Default Value	MSR Mask	Access Type	Description
<u>LOCK BIT</u>	Lock	31:31	0b	0b	<u>RW-KL</u>	This bit will lock all settings in this register for writes. The Leakage load line value in 10/256 mOhm increments. Values range from 0 mOhm to 10 mOhm.
<u>LEAK LOAD LINE R</u>	Static (Leakage) Load Line Impedance	23:16	36h	00h	<u>RW-L</u>	The default value of 36h corresponds to 2.1mOhm. The Idle load line value in 10/256 mOhm increments. Values range from 0 mOhm to 10 mOhm.
<u>IDLE LOAD LINE R</u>	Idle Load Line Impedance	15:8	36h	00h	<u>RW-L</u>	The default value of 36h corresponds to 2.1mOhm. The Dynamic load line value in 10/256 mOhm increments. Values range from 0 mOhm to 10 mOhm.
<u>DYNAMIC LOAD LINE R</u>	Dynamic Load Line Impedance (Load_Line)	7:0	36h	00h	<u>RW-L</u>	The default value of 36h corresponds to 2.1mOhm.

FIG. 4C



FIG. 4D

MSR Address:194h Access:Read/Write			
Bit	Scope	Default	Description
63:21	-	-	Reserved
20	Package	0	<b>OC_LOCK (RW/L)</b> When this bit is set bits 20, 7:0 in this register and all bits in TURBO_RATIO MSR 1ADh are locked. A reset is required to unlock any of these fields. Note: This bit is available only on Ivy Bridge processors with signature great than or equal to 000306A3h
19:17	Package	000	<b>Number of overclocking bins supported (RO)</b> This field indicates the maximum number of bins by which the part can be overclocked 000b-No overclocking support on the part 001b-1 bin overclocking supported 010b-2 bin overclocking supported 011b-3 bin overclocking supported 100b-4 bin overclocking supported 101b-5 bin overclocking supported 110b-6 bin overclocking supported 111b-Unlimited overclocking supported
16	Package	0	<b>FLEX_EN (R/O)</b> Flex Ratio Enabled. When this bit is set it means that the value in the FLEX_RATIO will be used. If this bit is clear, then value in the FLEX_RATIO field is ignored.
15:8	Package	0	<b>MAX_RATIO (R/W)</b> Desired Flex ratio. When enabled, the value in this field will set the maximum non-turbo ratio available. If this value is greater than the maximum ratio set by hardware, this field is ignored.
7:0	Package	0	<b>Over Clocking Extra Voltage (R/W)</b> Desired Voltage to be used for over clocking. The processor will use this value to extrapolate voltages for various Turbo frequencies when over clocking is enabled. The voltage value is specified in units of 1/256 Volts.




FIG. 6




FIG. 7




FIG. 8




FIG. 9

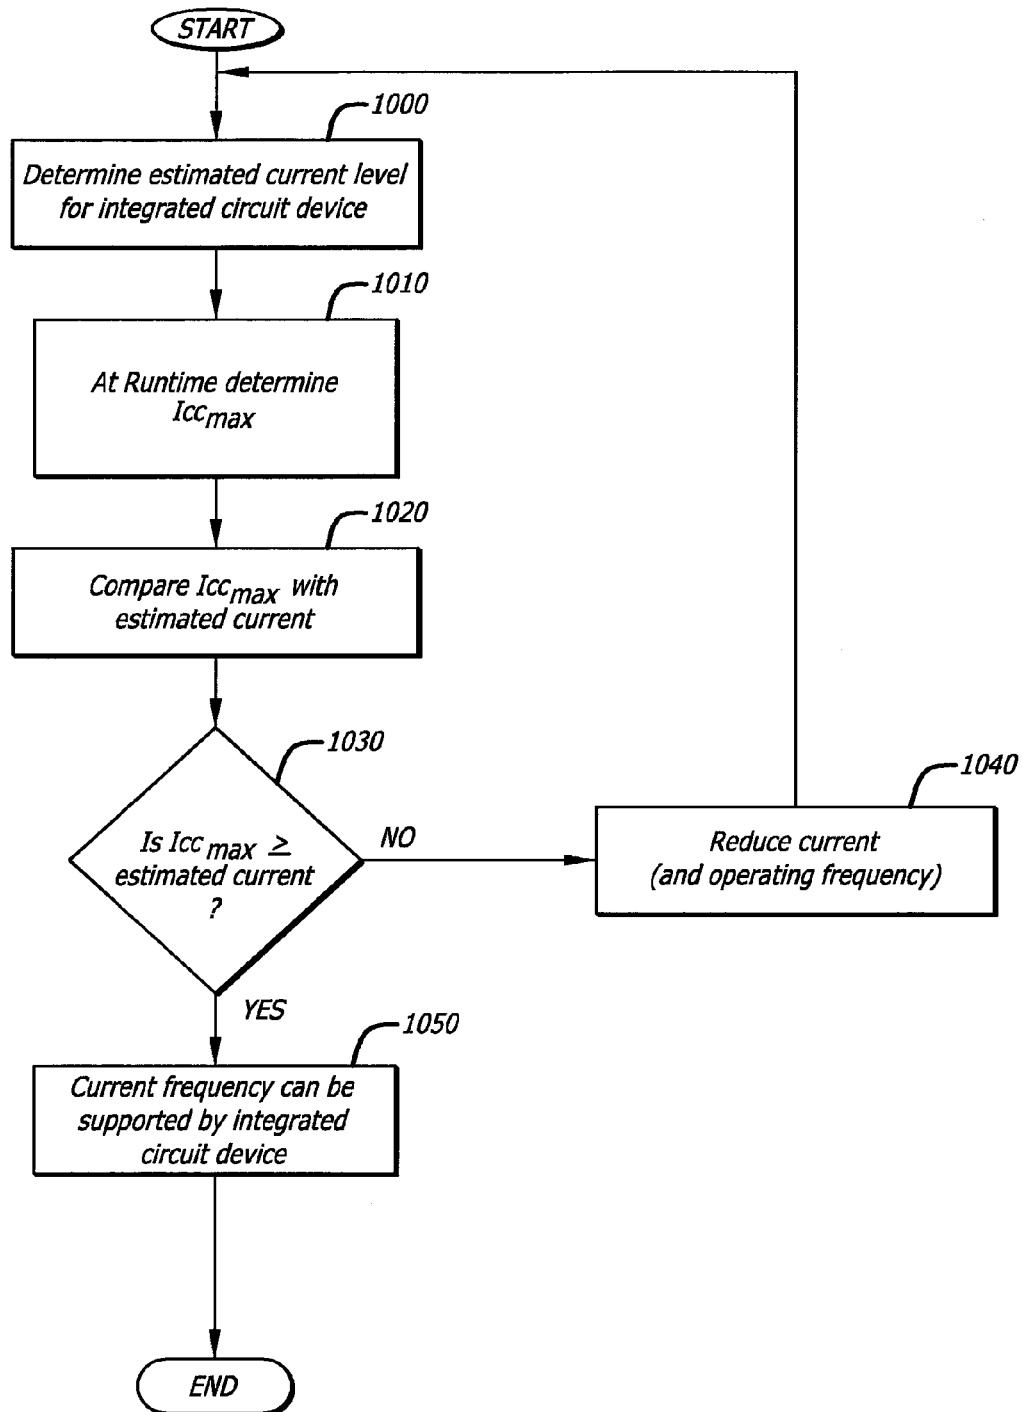


FIG. 10


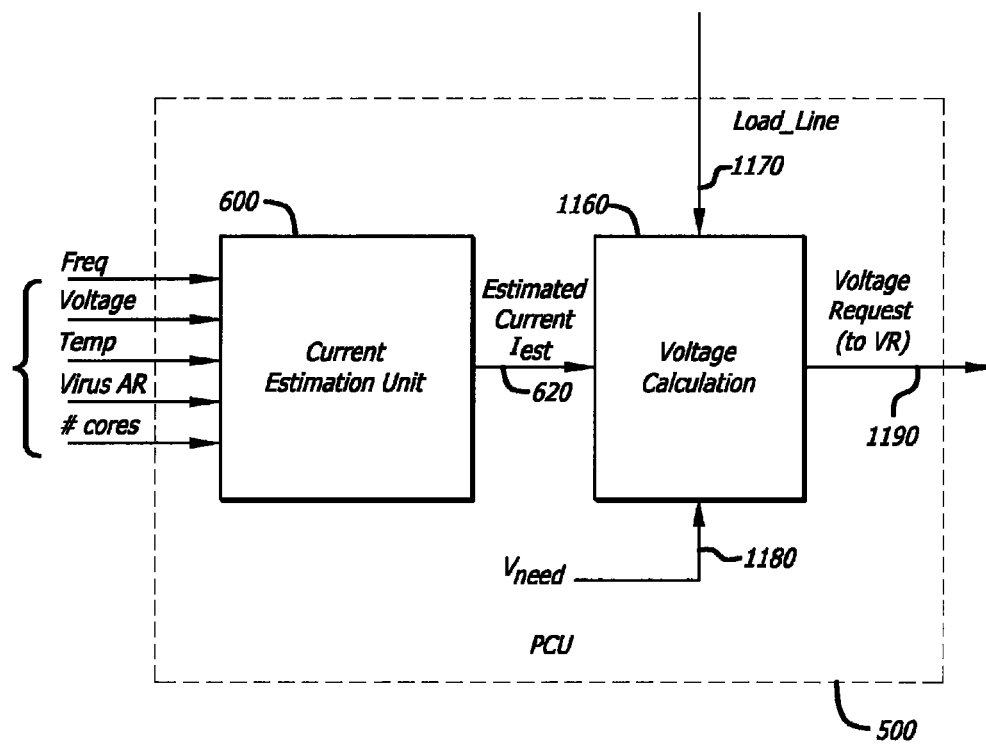
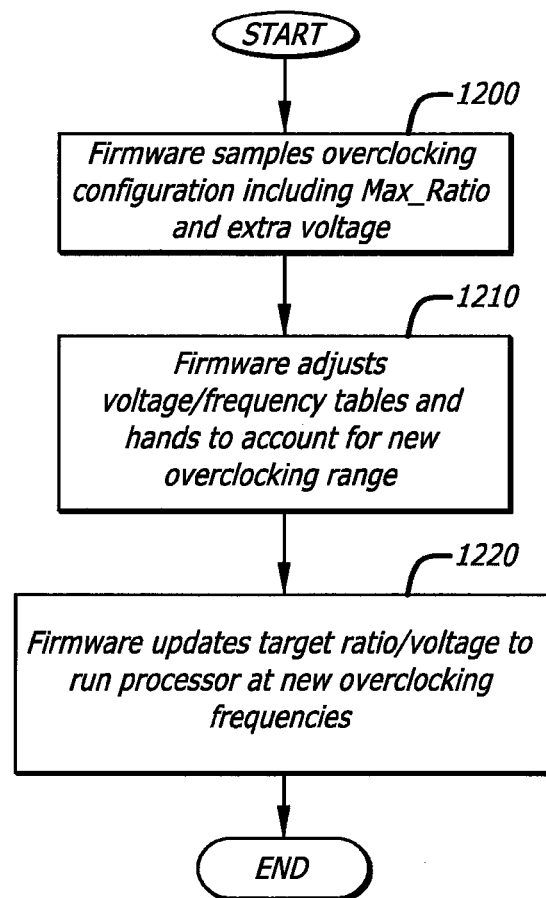



FIG. 11A

**FIG. 11B**

**FIG. 12**

**FIG. 13**






FIG. 14

1

# SYSTEM, METHOD AND APPARATUS FOR ENERGY EFFICIENCY AND ENERGY CONSERVATION BY CONFIGURING POWER MANAGEMENT PARAMETERS DURING RUN TIME

This application claims priority to Provisional Patent Application No. 61/635,817, entitled: "System, Method and Apparatus For Energy Efficiency And Energy Conservation By Configuring Power Management Parameters During Run Time" filed Apr. 19, 2012 and is hereby incorporated by reference in its entirety into this application.

## FIELD

Embodiments of the invention pertain to energy efficiency and energy conservation in integrated circuits, as well as code to execute thereon, and in particular but not exclusively, to run time programmability of various power management parameters.

## GENERAL BACKGROUND

Advances in semiconductor processing and logic design have permitted an increase in the amount of logic that may be present on integrated circuit devices. As a result, computer system configurations have evolved from a single or multiple integrated circuits in a system to multiple hardware threads, multiple cores, multiple devices, and/or complete systems on individual integrated circuits. Additionally, as the density of integrated circuits has grown, the power requirements for computing systems (from embedded systems to servers) have also escalated. Furthermore, software inefficiencies, and its requirements of hardware, have also caused an increase in computing device energy consumption. In fact, some studies indicate that computing devices consume a sizeable percentage of the entire electricity supply for a country, such as the United States of America. As a result, there is a vital need for energy efficiency and conservation associated with integrated circuits. These needs will increase as servers, desktop computers, notebooks, ultrabooks, tablets, mobile phones, processors, embedded systems, etc. become even more prevalent (from inclusion in the typical computer, automobiles, and televisions to biotechnology).

As general background, controlling power consumption in microprocessors and other integrated circuit devices has increased in importance, especially with the greater use of mobile devices. Some prior art techniques for managing processor power consumption have not adequately provided a dynamic scheme for setting various power management parameters relied upon by an integrated circuit device, such as a processor. The lack of a dynamic setting scheme for various power management parameters, other than the Thermal Design Power (TDP) parameter, not only lessens the actual power savings realized, but also restricts the ability of Original Equipment Manufacturers (OEMs) to design products that can temporarily operate outside specifications established by the processor manufacturer, such as Intel Corporation of Santa Clara, Calif.

## BRIEF DESCRIPTION OF THE DRAWINGS

The invention may best be understood by referring to the following description and accompanying drawings that are used to illustrate embodiments of the invention.

2

FIG. 1 is an exemplary block diagram of an electronic device implemented with an integrated circuit device with dynamic power management monitoring for performance adjustment.

FIG. 2 is a first exemplary block diagram of the system architecture implemented within the electronic device of FIG. 1 or another electronic device.

FIG. 3 is a second exemplary block diagram of the system architecture implemented within the electronic device of FIG. 1 or another electronic device.

FIG. 4A is a first exemplary block diagram of the packaged integrated circuit device with current monitoring as a single-core or multi-core processor having an integrated graphics and system agent.

FIGS. 4B-4D are illustrative embodiments of different software accessible registers adapted for storing different power management parameters.

FIG. 5 is an exemplary block diagram of a Power Control Unit (PCU) implemented within the system agent unit of the processor of FIG. 4A.

FIG. 6 is an exemplary embodiment of current estimation and computations to adjust the operating frequency of the processor of FIG. 4A in accordance with an embodiment of the invention.

FIG. 7 is a second exemplary block diagram of the packaged integrated circuit device with activity monitoring as a packaged multi-processor unit with at least one of the processors supporting power management monitoring.

FIG. 8 is a third exemplary block diagram of the packaged integrated circuit device with power management monitoring implemented on a circuit board.

FIG. 9 is an exemplary block diagram of an integrated circuit device with power management monitoring implemented within a blade server that is in communication with other blade servers.

FIG. 10 is an exemplary flowchart of the operations conducted by the integrated circuit device for dynamic monitoring and application of a maximum current ( $I_{cc_{max}}$ ) power management parameter.

FIG. 11A is an exemplary flowchart of the operations conducted by the integrated circuit device for dynamic compensation of load line drop between a voltage regulator and the processor.

FIG. 11B is an exemplary embodiment of adjustment of a voltage request during run time in accordance with an embodiment of the invention.

FIG. 12 is an exemplary flowchart of the operations conducted by the integrated circuit device for dynamic monitoring and application of a maximum ratio to control overclocking of the processor.

FIG. 13 shows a method that uses a lock bit.

FIG. 14 shows direct control of PCU constraints by application software.

## DETAILED DESCRIPTION

Herein, certain embodiments of the invention relate to an integrated circuit device adapted to dynamically monitor and adjust parameters relied upon for power management during run time of an electronic device implemented with the integrated circuit, namely after Basic Input Output System (BIOS) has completed and the operating system (OS) is in control of the electronic device. These power management parameters may include, but are not limited or restricted to (1) a parameter ( $I_{cc_{max}}$ ) representing a maximum current level that can be used by the integrated circuit device in a certain operating mode, (2) a parameter accounting for dynamic

changes in load line impedance (Load\_Line) as detected by logic within the electronic device, and (3) a parameter (Max\_Ratio) setting an overclocking condition optionally accompanied with an application of additional voltage. Through compliance with these dynamically adjustable power management parameters for a processor, for example, power conservation may be realized.

A corresponding method for dynamically adjusting power management parameters and an electronic device implementing such integrated circuit devices is also described herein.

In the following description, certain terminology is used to describe certain features of the invention. For example, the term “integrated circuit device” generally refers to any integrated circuit or collection of integrated circuits that are adapted to control its performance, and thus its power usage, using one or more power management parameters that are dynamically adjustable during run time. The power management parameters (e.g., maximum current level, maximum ratio for overclocking, etc.) are set through the use of storage elements that are software accessible, such as machine specific registers (MSRs) for example. For instance, the maximum current level may be set through the use of PP0\_CURRENT\_CONFIG and PP1\_CURRENT\_CONFIG registers described below. Examples of an integrated circuit device may include, but are not limited or restricted to a processor (e.g. a single or multi-core microprocessor, a digital signal processor “DSP”, or any special-purpose processor such as a network processor, co-processor, graphics processor, embedded processor), a microcontroller, an application specific integrated circuit (ASIC), a memory controller, an input/output (I/O) controller, or the like.

Furthermore, the term “logic” constitutes hardware and/or software. As hardware, logic may include processing circuitry (e.g., a controller, processor, an application specific integrated circuit, etc.), semiconductor memory, combinatorial logic, or the like. As software, the logic may be one or more software modules, such as executable code in the form of an executable application, an application programming interface (API), a subroutine, a function, a procedure, an object method/implementation, an applet, a servlet, a routine, a source code, an object code, firmware, a shared library/dynamic load library, or one or more instructions.

It is contemplated that these software modules may be stored in any type of suitable non-transitory storage medium or transitory computer-readable transmission medium. Examples of non-transitory storage medium may include, but are not limited or restricted to a programmable circuit; a semiconductor memory such as a volatile memory such as random access memory “RAM,” or non-volatile memory such as read-only memory, power-backed RAM, flash memory, phase-change memory or the like; a hard disk drive; an optical disc drive; or any connector for receiving a portable memory device such as a Universal Serial Bus “USB” flash drive. Examples of transitory storage medium may include, but are not limited or restricted to electrical, optical, acoustical or other form of propagated signals such as carrier waves, infrared signals, and digital signals.

The term “interconnect” is broadly defined as a logical or physical communication path for information. This interconnect may be established using any communication medium such as a wired physical medium (e.g., a bus, one or more electrical wires, trace, cable, etc.) or a wireless medium (e.g., air in combination with wireless signaling technology).

Lastly, the terms “or” and “and/or” as used herein are to be interpreted as inclusive or meaning any one or any combination. Therefore, the phrases “A, B or C” and “A, B and/or C” mean any of the following: A; B; C; A and B; A and C; B and

C; A, B and C. An exception to this definition will occur only when a combination of elements, functions, steps or acts are in some way inherently mutually exclusive.

#### I. System Architecture

Referring now to FIG. 1, an exemplary block diagram of an electronic device 100, which is implemented with one or more integrated circuit devices with run time adjustable power management parameters, is shown. Herein, electronic device 100 is realized, for example, as a notebook-type personal computer. However, it is contemplated that electronic device 100 may be a desktop computer, a television, a portable device, or an embedded application. Examples of a “portable device” may include, but is not limited or restricted to a cellular telephone, any portable computer including a tablet computer, an Internet Protocol (IP) device, a digital camera, a personal digital assistant (PDA), a video game console, a portable music player, or a digital camera. An example of an “embedded application” typically includes a microcontroller, a digital signal processor (DSP), a system-on-a-chip, network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, or any other system that can perform the functions and operations taught below.

As will become readily apparent in the description below, the embodiments of methods, electronic devices, and systems described herein (whether in reference to hardware, firmware, software, or a combination thereof) are vital to a ‘green technology’ future, such as for power conservation and energy efficiency in products that encompass a large portion of the US economy.

As shown in FIG. 1, electronic device 100 includes a housing 110 and a display unit 120. According to this embodiment of the invention, display unit 120 includes a liquid crystal display (LCD) 130 which is built into display unit 120. According to one embodiment of the invention, display unit 120 may be rotationally coupled to housing 110 so as to rotate between an open position where a top surface 112 of housing 110 is exposed, and a closed position where top surface 112 of housing 110 is covered. According to another embodiment of the invention, display unit 120 may be integrated into housing 110.

Referring still to FIG. 1, housing 110 may be configured as a box-shaped housing. According to one embodiment of the invention, an input device 140 is disposed on top surface 112 of housing 110. As shown, input device 140 may be implemented as a keyboard 142 and/or a touch pad 144. Although not shown, input device 140 may be touch-screen display 130 that is integrated into housing 110, or input device 140 may be a remote controller if electronic device 100 is a television.

Other features include a power button 150 for powering on/off electronic device and speakers 160<sub>1</sub> and 160<sub>2</sub> disposed on top surface 112 of housing 110. At a side surface 114 of housing 110 is provided a connector 170 for downloading or uploading information (and/or applying power). According to one embodiment, connector 170 is a Universal Serial Bus (USB) connector although another type of connector may be used.

As an optional feature, another side surface of electronic device 100 may be provided with high-definition multimedia interface (HDMI) terminal which support the HDMI standard, a DVI terminal or an RGB terminal (not shown). The HDMI terminal and DVI terminal are used in order to receive or output digital video signals with an external device.

Referring now to FIG. 2, a first exemplary block diagram of the system architecture implemented within electronic device 100 of FIG. 1 is shown. Herein, electronic device 100 comprises one or more processors 200 and 210. Processor 210 is

shown in dashed lines as an optional feature as electronic device **100** may be adapted with a single processor as described below. Any additional processors, such as processor **210**, may have the same or different architecture as processor **200** or may be an element with processing functionality such as an accelerator, field programmable gate array (FPGA), DSP or the like.

Herein, processor **200** comprises an integrated memory controller (not shown), and thus, is coupled to memory **220** (e.g., non-volatile or volatile memory such as a random access memory “RAM”). Furthermore, processor **200** is coupled to a chipset **230** (e.g., Platform Control Hub “PCH”) which may be adapted to control interaction between processor(s) **200** and **210** and memory **220** and incorporates functionality for communicating with a display device **240** (e.g., integrated LCD) and peripheral devices **250** (e.g., input device **140** of FIG. 1, wired or wireless modem, etc.). Of course, it is contemplated that processor **200** may be adapted with a graphics controller (not shown) so that display device **240** may be coupled to processor **200** via a Peripheral Component Interconnect Express (PCI-e) port **205**.

Referring now to FIG. 3, a second exemplary block diagram of the system architecture implemented within electronic device **100** of FIG. 1 is shown. Herein, electronic system **100** is a point-to-point interconnect system, and includes first processor **310** and second processor **320** coupled via a point-to-point (P-P) interconnect **330**. As shown, processors **310** and/or **320** may be some version of processors **200** and/or **210** of FIG. 2, or alternatively, processor **310** and/or **320** may be an element other than a processor such as an accelerator or FPGA.

First processor **310** may further include an integrated memory controller hub (IMC) **340** and P-P circuits **350** and **352**. Similarly, second processor **320** may include an IMC **342** and P-P circuits **354** and **356**. Processors **310** and **320** may exchange data via a point-to-point (P-P) interface **360** using P-P circuits **352** and **354**. As further shown in FIG. 3, IMC **340** and IMC **342** couple processors **310** and **320** to their respective memories, namely memory **370** and memory **372**, which may be portions of main memory locally attached to respective processors **310** and **320**.

Processors **310** and **320** may each exchange data with a chipset **380** via interfaces **390** and **392** using P-P circuits **350**, **382**, **356** and **384**. Chipset **380** may be coupled to a first bus **395** via an interface **386**. In one embodiment, first bus **395** may be a Peripheral Component Interconnect Express (PCI-e) bus or another third generation I/O interconnect bus, although the scope of the present invention is not so limited.

Referring to FIG. 4A, a first exemplary block diagram of an integrated circuit device **400** supporting a dynamic power management parameter adjustment scheme is shown. Herein, integrated circuit device **400** constitutes multi-core processor **200**, which is partitioned to include a plurality of power planes **420**, **440** and **460**. Power plane **440** comprises a first compute engine **405** that includes multiple ( $N \geq 2$ ) processor cores **410**<sub>1</sub>-**410**<sub>N</sub>. However, it is contemplated that first compute engine **405** may include a single processor core **410**<sub>1</sub> and/or elements with processing functionality other than a processor core.

Herein, a control unit **430**, referred to herein the “system agent” (SA), is implemented on first power plane **420**. System agent **430** is responsible for adjusting frequency and voltage applied to power planes **440** and **460**. Moreover, system agent **430** is responsible for adjusting the power management parameters as described below.

As further shown in FIG. 4A, second power plane **440** includes processor cores **410**<sub>1</sub>-**410**<sub>N</sub> and an on-chip memory

architecture **445** coupled thereto. On-chip memory architecture **445** includes a plurality of software-accessible storage elements **450** (e.g., machine specific registers “MSRs”) and a plurality of on-chip memories **455**<sub>1</sub>-**455**<sub>M</sub> ( $M \geq 1$ ). On-chip memories **455**<sub>1</sub>-**455**<sub>M</sub> may be last-level caches (LLCs) each corresponding to one of the processor cores **410**<sub>1</sub>-**410**<sub>N</sub> ( $M=N$ ).

Herein, MSRs **450** are on-chip storage elements that are programmed to store values that represent limits, impedances, multipliers, or other information that is accessible by SA **430** for power management of integrated circuit device **400**. MSRs **450** may be implemented on second power plane **440** as shown. However, it is contemplated that MSRs **450** may be implemented on power planes **420** or **460** such as part of a Power Control Unit (PCU) within SA **430**.

As shown, a first MSR **451** may include a value ( $I_{cc,max}$ ) **480** that provides SA **430** with the maximum amount of instantaneous current for a particular power plane as shown in FIG. 4B. In particular, when CURRENT\_LIMIT\_LOCK bit **481** is not set, CURRENT\_LIMIT bits **482** may be adjusted during run time to represent the maximum current limit that can be realized by processor **400** of FIG. 4A, which will adjust its operating conditions in order to meet the new limit. Such changes may be due to changes in the power delivery for the electronic device, such as the addition of another device (e.g., Universal Serial Bus “USB” or Firewire) which requires reallocation of device power delivery resources.

As shown in FIG. 4C, a second MSR **452** may include a value **485** that provides SA **430** with impedance for an interconnect (load line) between a power source (e.g., voltage regulator) and a compute engine (e.g. first compute engine **405**). This value (Load\_Line) **485** may dynamically change due to operational variations detected by electronic device **100**. Examples of these variations may include, for example, a change in the power source (e.g. switch from an alternating current “AC” source to direct current “DC” source or vice versa), a change in operating state of electronic device **100** (e.g. from normal operating state to sleep state), a change in the number of active threads being processed, or the like. Load\_Line **485**, along with the known current being supplied are used by SA **430**, namely firmware (P-code), is used to adjust the amount of voltage provided to second power plane **420** (and indirectly to compute engine **405**) to offset a voltage drop caused by impedance from the interconnect between processor **400** and its power source. The impedance may be changed at runtime and the processor will adjust its operating conditions in order to comprehend the new value.

As shown in FIG. 4D, a third MSR **453** is adapted to store Max\_Ratio **490**, which provides PCU **500** (described below) with a maximum ratio used for overclocking. More specifically, when conditions are deemed satisfactory to exceed the guaranteed TDP frequency, Max\_Ratio **490** is used as a new multiplier for a reference clock frequency to provide an increased clock frequency for at least one of processor cores **410**<sub>1</sub>-**410**<sub>N</sub>. As overclocking is the process of operating the compute engine at a higher clock rate (more clock cycles per second) than it was designed for or was specified by the manufacturer, third MSR **453** is further adapted to, by setting OVER\_CLOCKING\_EXTRA\_VOLTAGE bits **495**, signify an amount of extra voltage to be added in response to the increased operating frequency. Processor **400** updates its voltage and frequency tables in accordance with Max\_Ratio **490**.

Referring back to FIG. 4A, third power plane **460** includes a second compute engine, namely graphics logic **470**. Control of third power plane **460** is the same as control of second power plane **440**. Third power plane **460** is controlled independently

from second power plane 440. For instance, the  $I_{cc,max}$  for second power plane 440 may be different from  $I_{cc,max}$  for third power plane 460. Hence, MSRs 450 may include additional power management constraints for components implemented on different power planes.

In accordance with one embodiment of the invention, the adjustment of the power management parameters is managed by specific logic within system agent 430, namely a Power Control Unit (PCU) 500. A hybrid of hardware and firmware, PCU 500 is effectively a controller that manages all of the power management associated with integrated circuit device 400. Of course, in lieu of PCU 500, power management may be accomplished through an embedded controller within integrated circuit device 400, an off-die controller (e.g., controller within the same multi-chip package, on the same circuit board, etc.), or other types of logic.

## II. Power Management of the Integrated Circuit Device

Referring now to FIG. 5, one embodiment of system agent 430 features PCU 500 is shown. PCU 500 includes a micro-controller 510 that runs firmware 520 (P-code) for controlling the adjustment of various power management parameters in efforts to effect operations on the different power planes 440 and 460. In other words, according to one embodiment of the invention, P-code 520, when executed, is adapted to monitor operational components used for estimating and managing current usage and, where appropriate, permitting or precluding such current levels. This current adjustment effectively adjusts the operating frequency and power usage by one or more of processor cores 410<sub>1</sub>-410<sub>N</sub> implemented on power plane 440.

PCU 500 further includes hardware state machines 530 that control the transitioning in frequency (and voltage) for power planes 440 and 460.

According to one embodiment of the invention, as shown in FIGS. 4A and 5, each of processor cores 410<sub>1</sub>-410<sub>N</sub> is associated with assigned power parameters. Based on the current work level that is needed, PCU 500 is adapted to control processor cores 410<sub>1</sub>-410<sub>N</sub> and, based on the estimated current and predetermined current limits, allows or precludes such current usage. This effectively controls the operating frequency of processor cores 410<sub>1</sub>, . . . , and/or 410<sub>N</sub>. Such control is performed based on parameters as set forth in FIG. 6.

Referring now to FIG. 6, an exemplary embodiment of current estimation and computations performed by PCU 500 of FIG. 5 in controlling performance of integrated circuit device 400 of FIG. 4A is shown. Herein, PCU 500 of FIGS. 4A and 5 may need to alter, during run time, maximum current usage based on operational variations detected by the electronic device. The need for such alteration may be determined, for example, by comparing a maximum current threshold to an estimated current computed using parameters. These parameters may include, but are not limited or restricted any or all of the following: frequency, voltage, temperature, power virus (e.g., 100% app ratio being the highest power level by integrated circuit device 400), the number of cores, reference leakage and/or leakage scaling.

More specifically, current estimation unit 600 of PCU 500 actively or passively monitors parameters 610 that are used to estimate an amount of current utilized by the integrated circuit device. For example, according to one embodiment of the invention, values as measured by sensors or contained within MSRs are input into current estimation unit 600.

According to this embodiment of the invention, as illustrated in FIG. 6, current estimation unit 600 performs an arithmetic operation on parameters 610 in order to produce estimated current ( $I_{est}$ ) 620. Herein, the estimated current is

computed by PCU 500, and in particular P-code 520, as a function of (i) frequency, (ii) voltage as a function of frequency, (iii) current temperature of the integrated circuit device, (iv) power virus, (v) number of processor cores within integrated circuit device 400, and perhaps reference leakage and leakage scaling parameters (not shown) or the like. Of course, it is contemplated that  $I_{est}$  620 may be computed by logic other than P-code 520 and provided as input to P-code 520.

Regardless, estimated current ( $I_{est}$ ) 620 is input into a current limiting unit 630. Current limiting unit 630 is adapted to compare  $I_{est}$  620 to a current limit 640, which is accessible by P-code 520 from a MSR. As an illustrative example, a storage element 650 may be adapted to store a value that represents a maximum current. For instance, PP0\_CURRENT\_CONFIG MSR may be adapted to identify a maximum instantaneous current ( $I_{cc,max}$ ) for the power rail utilized by processor core(s) 410<sub>1</sub>, . . . , and/or 410<sub>N</sub> within power plane 440.  $I_{cc,max}$  is represented by bits [12:0] of PP0\_CURRENT\_CONFIG MSR as shown in FIG. 4B. Similarly, another MSR (PP1\_CURRENT\_CONFIG) may be adapted to identify a maximum current ( $I_{cc,max}$ ) for the power rail utilized by graphics logic 470 within power plane 460.

Thereafter, during run time, Current Limiting unit 630 compares the estimated current to  $I_{cc,max}$  in order to determine if the estimated current is equal to or falls below  $I_{cc,max}$ . If the estimated current is equal to or falls below  $I_{cc,max}$ , then the proposed current usage is sustainable and the operating frequency corresponding to this current usage level is permitted. However, if the estimated current is greater than  $I_{cc,max}$ , PCU 500 reduces the current, which reduces the operating frequency and corresponding voltage (P-State) of a particular compute engine, such as one of the processor cores 410<sub>1</sub>-410<sub>N</sub> of FIG. 4A. Of course, it is contemplated that such computations may be conducted based on average, median or worst case instantaneous current usage with short time window, or the like.

Referring now to FIG. 7, a second exemplary block diagram of an integrated circuit device 700 as a packaged multi-processor unit with at least one of the processors supporting a power management parameter adjustment scheme is shown. Herein, packaged integrated circuit device 700 includes a package 710 partially or fully encapsulating a substrate 720. Substrate 720 comprises a controller 730 that is adapted to monitor and limit current usage based on current constraints applicable to packaged integrated circuit device 700 and/or multiple integrated circuit devices 740<sub>1</sub>-740<sub>P</sub> ( $P \geq 2$ ) on substrate 720. Hence, controller 730 performs the above-described operations of the PCU implemented in accordance with the integrated circuit (die) architecture shown in FIG. 4A.

Referring now to FIG. 8, a third exemplary block diagram of the packaged integrated circuit device implemented on a circuit board 800 is shown. Herein, a packaged integrated circuit device 810 is mounted on circuit board 800 and is adapted as a controller for power management during run time. For instance, integrated circuit device 810 may be adapted to effectively monitor and limit currents applied to different portions of integrated circuit device 810 or other integrated circuit devices 820<sub>1</sub>-820<sub>P</sub> ( $P \geq 1$ ) on circuit board 800. Hence, controller 810 performs the above-described operations of the PCU implemented in accordance with the integrated circuit (die) architecture shown in FIG. 4A.

Referring to FIG. 9, an exemplary block diagram of the electronic device implemented with an integrated circuit device 900 implemented with a blade server 910 for activity monitoring is shown. Herein, packaged integrated circuit

device **900** is adapted as a controller to monitor and limit frequency and voltage levels for one or more different blade servers **920** other than blade server **910**. Hence, controller **900** performs the above-described operations of the PCU implemented in accordance with the integrated circuit (die) architecture shown in FIG. 4A.

Referring now to FIG. 10, an exemplary flowchart of the operations by firmware within an integrated circuit device for dynamically monitoring and adjusting power management parameters for power conservation is shown. First, during run-time, firmware determines an estimated current level based on a plurality of deterministic parameters (block **1000**). For instance, as an illustrative embodiment of the invention, the estimated current level may be computed using voltage, frequency, temperature, power virus and number of cores. Of course, it is contemplated that the estimated current level may be determined by logic other than firmware and/or a different estimation scheme may be used that includes additional parameters and/or exclude some of the listed parameters. Regardless, the estimated current level for the integrated circuit device may be calculated periodically, randomly or in response to a triggering event during run time.

During run time of the integrated circuit device, perhaps in response to a change in condition of the integrated circuit device **400** of FIG. 4A or the electronic device **100** of FIG. 1,  $I_{cc,max}$  is obtained from a designated storage element such as a software accessible register (block **1010**). This  $I_{cc,max}$  may vary depending on which power plane of the integrated circuit (or area on the substrate or circuit board). For instance,  $I_{cc,max}$  for the processor power plane may be obtained from PP0\_CURRENT\_CONFIG MSR while  $I_{cc,max}$  for the graphics power plane may be obtained from PP1\_CURRENT\_CONFIG MSR, where these MSRs may store different  $I_{cc,max}$  values.  $I_{cc,max}$  may be adjusted during run time by the electronic device and will be re-sampled and considered by the firmware.

Thereafter, the estimated current level is compared to  $I_{cc,max}$  (block **1020**). Where  $I_{cc,max}$  is associated with amperage greater than the estimated current level, the current level can be supported by the integrated circuit (blocks **1030** and **1040**). Thus, no current adjustment is needed at this time and the integrated circuit device can operate at its existing operating frequency. Otherwise, the current is reduced which effectively reduces the operating frequency of the integrated circuit (blocks **1030** and **1050**).

Referring to FIG. 11, an exemplary flowchart of the operations by firmware within an integrated circuit device for dynamically monitoring and adjusting power management parameters, such as a parameter corresponding to estimated impedance of the load line between the integrated circuit device and a power source (e.g., voltage regulator), is shown. This parameter, hereinafter referred to as "Load\_Line," can be used to dynamically alter the supply voltage from the power source to the integrated circuit device in order to more accurately account for voltage loss realized over the load line. Such accuracy enables power conservation as the worst case load line impedance does not need to be used for computing the amount of voltage applied to the integrated circuit device from the power source.

First, according to one embodiment of the invention, an impedance of the load line is calculated (block **1100**). For instance, as an illustrative embodiment of the invention, this impedance may be computed by circuitry within the electronic device (e.g., applying a constant current and measuring the voltage at each end of the load line). Once this impedance is determined, a value (Load\_Line), which represents the estimated load line impedance, is stored in an accessible

register (block **1110**). For instance, Load\_Line for the processor power plane may be stored within a particular storage element (e.g. bits [7:0] of PRIMARY\_PLANE\_LLRCONFIG\_CONTROL MSR). As an example, Load\_Line may represent an impedance ranging anywhere from 0-10 milliohms ( $m\Omega$ ). Likewise, Load\_Line for the graphics power plane may be stored in another software-accessible storage element (e.g. bits [7:0] of SECONDARY\_PLANE\_LLRCONFIG\_CONTROL MSR) which represents an impedance ranging anywhere from 0-10  $m\Omega$ .

Next, as shown, in response to an event during run time (e.g., action, received signaling, elapsed time, etc.) which may occur periodically or randomly, firmware accesses Load\_Line (blocks **1120** and **1130**). Thereafter, as shown in block **1140**, Load\_Line is used to determine the amount of voltage drop caused by the load line between the integrated circuit device (e.g., processor **400** of FIG. 4A) and its power source. As a result, firmware uses this information, in combination with the known current supplied to the processor, to request additional voltage from the power source that matches the voltage drop along the load line (block **1150**). This avoids usage of worst case voltage losses which normally causes the delivery of higher voltage to the processor than needed.

In general terms, the processor reads load line impedance (Load\_Line) and using Load\_Line to calculate a voltage drop over the load line. Thereafter, the amount of requested voltage compensates for the calculated voltage drop.

As shown in FIG. 11B, current estimation unit **600** of PCU **500** actively or passively monitors parameters **610** that are used to estimate an amount of current utilized by the integrated circuit device. For example, according to one embodiment of the invention, these values (measured by sensors or contained within MSRs) are input into current estimation unit **600**.

According to this embodiment of the invention, as illustrated in FIG. 11B, current estimation unit **600** performs an arithmetic operation on parameters **610** in order to produce estimated current ( $I_{est}$ ) **620** as previously described. Estimated current ( $I_{est}$ ) **620** is input into a voltage calculation circuit **1160**, which takes the voltage needed by the processor ( $V_{need}$ ) **1180**, estimated current ( $I_{est}$ ) **620** and Load\_Line **1170** from the electronic device to produce a voltage request **1190**. Herein, the load line impedance is read and readjusting voltage request at runtime.

Referring now to FIG. 12, an exemplary flowchart of the operations by firmware within an integrated circuit device for dynamically monitoring and adjusting power management parameters, namely the maximum ratio that is relied upon for setting the operating frequency of the integrated circuit device, is shown. In particular, the maximum ratio is used to set an overlocking condition.

In general, for run-time overlocking, the firmware (e.g., P-code) samples overlocking configuration including maximum ratios and extra voltage (block **1200**). Thereafter, the firmware adjusts voltage/frequency tables and limits to account for a new overlocking range (block **1210**). The firmware then updates the target ratio/voltage to run processor at new overlocking frequencies (block **1220**).

More specifically, when permitted, firmware samples the Max\_Ratio **490** of FIG. 4D and uses this ratio as a multiplier for a reference clock frequency. This causes overlocking by increasing the operating frequency for at least one of processor cores **410<sub>1</sub>-410<sub>N</sub>** beyond the established TDP frequency and beyond a maximum turbo frequency spec'ed by the manufacturer of the integrated circuit device. Furthermore, firmware samples the OVER\_CLOCKING\_EXTRA\_VOLT-

11

AGE bits **495** of FIG. **4D** and updates its voltage and frequency tables to account for the increased voltage and frequency.

The aforementioned discussion provided various implementation details for a mechanism that permits changes during run-time to various constraints (e.g., current limit, load line and overlocking) that are utilized by a power control unit (PCU) that determines appropriate clock frequency, clock gating and supply voltage values in view of the constraints.

FIGS. **4b**, **4c** and **4d** disclose a respective lock bit for each of the current threshold, load line and overlocking constraints. As alluded to in the descriptions of these lock bits in FIGS. **4b**, **4c** and **4d**, each lock bit, once set, is designed to prevent any further changes to its respective constraint until a reset occurs. Each lock bit can be used, for instance, to disable software's ability to change a constraint setting once true run time commences.

FIG. **13** shows an example. Initially, the computing system processor experiences a reset, for example, as part of the system's normal initial power on sequence **1301**. After coming out of reset, during the initial power on sequence, the processor's corresponding processing cores run out of BIOS firmware.

Upon power on/reset the lock bit **1302** is cleared and BIOS firmware subsequently looks to see if the configuration settings for the system permit run time changes to PCU constraints **1303**. Here, as the ability to change power management constraint settings during run time may result in settings that exceed a processing core's "guarantees" (that is, the settings cause the processor to behave unreliably), some system designers may wish to disable the ability to change power constraints during run time. Such system designers may therefore, for example, set a field in firmware that indicates that the ability to change constraints during run time is to be disabled.

BIOS firmware looks to this field **1303**, and, if the field indicates that run time setting changes are to be disabled **1304**, BIOS sets the appropriate lock bit(s) in MSR register space **1305**. Once a lock bit is set, changes to the power control unit constraint are not permitted until after another reset event occurs. This prevents software to change a constraint setting during runtime.

As observed with the multiple respective lock bits of FIGS. **4b**, **4c** and **4d**, according to one embodiment, each constraint parameter may have its own respective lock bit. In this case, in order to completely disable run time change capability, BIOS firmware sets the lock bit for each constraint parameter (e.g., lock bits for each of current threshold, load line and overlocking would have to be set by BIOS firmware). Note that having a respective lock bit for each constraint provides system manufacturers with a fine grained ability to impose different flavors of restraint to the run time PCU constraint change technology. For example, a system manufacturer may permit current threshold constraint setting changes during run time but not overlocking changes by setting the over clock lock bit of FIG. **4D** but not setting the current limit lock bit of FIG. **4B**.

In another embodiment (not shown) a global lock bit is provided in MSR register space that has the same effect as setting a lock bit for each constraint parameter that can be changed (e.g., setting each of the current limit, load line and overlocking lock bits). As such, setting the global lock bit prevents software from changing any PCU constraint during run time. Some implementations may choose to have only a global lock bit without fine grained lock bits for each settable constraint, or, have the global lock bit along with the fine grained lock bits in MSR space.

12

FIG. **14** shows a software platform for enabling a user and/or application software to set a PCU constraint during runtime. As observed in FIG. **14** a device driver **1401** having specialized code for writing to the appropriate MSR register space **1402** is embedded in the operating system (OS) **1403**. Here, the OS may be an instance of an operating system running "on top of" a virtual machine monitor (VMM). The term operating system (OS) is used herein to refer to an operating system instance and/or a VMM. The device driver, for example, has knowledge of the specific one or more addresses within MSR register space where values for the various PCU constraints (e.g., current limit, load line and overlocking) are kept, and, has permission to access MSR register space **1402** and write new constraint values accordingly. In an embodiment, MSR register space **1402** is accessed through the use of memory mapped I/O (MMIO) techniques (i.e., memory is used as a conduit to exchange information between the MSR registers and the device driver).

The device driver **1401** also has an interface **1404** that is made available to application software **1405** that runs on the OS **1403**. The interface **1404**, in an embodiment, provides functions that the application software **1405** can call on to control the PCU constraints explicitly. For example, in an embodiment where the constraints include current limit, load line, over clocking and over clocking voltage, the interface supports: 1) a read current limit constraint function; 2) a read load line constraint function; 3) a read over clocking constraint function; 4) a read over clocking voltage function; 5) a write current limit constraint function; 6) a write load line constraint function; 7) a write over clocking constraint function; 8) a write over clocking voltage function.

In this example, application software **1405** can call eight different functions through interface **1404** each of which cause, as a direct consequence of the call, the corresponding MSR register space to be accessed. When write functions are called, application software **1405**, rather than OS software, controls the run time changes made to the PCU constraints of the system.

According to one approach, application software **1405** include or is associated with a Graphical User Interface (GUI) **1406** that permits a user to explicitly specify a constraint. That is, for example, the user of the system can specify a particular overlocking ratio and/or a particular over clocking supply voltage and have the new settings take effect during runtime without OS shut down, CPU reset, reboot, etc. Moreover, these settings may exceed the underlying processor core's or cores' guaranteed ratings. For certain applications, such as gaming applications, the fall off in processor reliability in exchange for the increase in performance is deemed acceptable to permit operation beyond a processor's guaranteed ratings.

The device also permits other new usage cases, such as, power management control at the application layer. That is, user application software (instead of embedded OS functions) controls system power consumption through interface **1404**. Here, application software **1405** may correspond to application software that is designed to provide a GUI **1406** that permits the user to specify specific power consumption targets and/or specific constraint values for the PCU as a function of system usage. For example, the user can specify a first consumption target or set of constraints when the system is running a first combination of applications (or a specific application) and a second set of targets/constraints when the system is running a second combination of applications (or specific application).

13

In the preceding examples, explicit reference was made to current limit, load line, over clocking and voltage over clocking PCU constraints. It should be understood that other embodiments may also include other kinds of constraints such as any one or more of: maxim turbo ratio per active processing core; graphics core ratio; graphics core voltage, specific power consumption limits, etc.).

While the invention has been described in terms of several embodiments, the invention should not be limited to only those embodiments described, but can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of limiting.

What is claimed is:

1. A method comprising:  
providing an interface of a processor to application software during runtime that permits at least one of a user and said application software to set a power management constraint for a power control unit of the processor during runtime of said processor without a reboot of the processor, wherein the power management constraint comprises a plurality of constraint parameters and the method further comprises clearing a respective lock bit for each of the plurality of constraint parameters on a reset event of the processor, wherein each lock bit disables the interface from changing a respective constraint parameter when set.
2. The method of claim 1, wherein the interface of the processor permits the user and the application software to set the power management constraint during runtime of the processor without the reboot of the processor.
3. The method of claim 1, wherein the power management constraint comprises a first set of constraints for when the processor is running a first application and a second set of constraints for when the processor is running a second application specified by the user.
4. The method of claim 1, wherein the power management constraint comprises a first consumption target for when the processor is running a first application and a second consumption target for when the processor is running a second application specified by the user.
5. The method of claim 1, wherein the power management constraint comprises a first set of constraints for when the processor is running a first combination of applications and a second set of constraints for when the processor is running a second combination of applications specified by the user.
6. The method of claim 1, wherein the power management constraint comprises a first consumption target for when the processor is running a first combination of applications and a second consumption target for when the processor is running a second combination of applications specified by the user.
7. The method of claim 1, further comprising clearing a lock bit on a reset event of the processor, wherein the lock bit disables the interface when set.
8. The method of claim 1, wherein the plurality of constraint parameters are overclocking constraint parameters.
9. A processor comprising:  
a power control unit; and  
an interface to application software during runtime that permits at least one of a user and said application software to set a power management constraint for the power control unit during runtime of said processor without a reboot of the processor, wherein the power management constraint comprises a plurality of constraint parameters and the processor further comprises at least one register to store a respective lock bit for each of the plurality of

14

constraint parameters to disable the interface from changing a respective constraint parameter when set.

10. The processor of claim 9, wherein the interface of the processor is to permit the user and the application software to set the power management constraint during runtime of the processor without the reboot of the processor.

11. The processor of claim 9, wherein the power management constraint comprises a first set of constraints for when the processor is to run a first application and a second set of constraints for when the processor is to run a second application to be specified by the user.

12. The processor of claim 9, wherein the power management constraint comprises a first consumption target for when the processor is to run a first application and a second consumption target for when the processor is to run a second application to be specified by the user.

13. The processor of claim 9, wherein the power management constraint comprises a first set of constraints for when the processor is to run a first combination of applications and a second set of constraints for when the processor is to run a second combination of applications to be specified by the user.

14. The processor of claim 9, wherein the power management constraint comprises a first consumption target for when the processor is to run a first combination of applications and a second consumption target for when the processor is to run a second combination of applications to be specified by the user.

15. The processor of claim 9, further comprising a register to store a lock bit to disable the interface when set.

16. The processor of claim 9, wherein the plurality of constraint parameters are overclocking constraint parameters.

17. A non-transitory computer-readable storage medium storing code that when executed by a computer causes the computer to perform a method comprising:

providing an interface of a processor to application software during runtime that permits at least one of a user and said application software to set a power management constraint for a power control unit of the processor during runtime of said processor without a reboot of the processor, wherein the power management constraint comprises a plurality of constraint parameters and the method further comprises clearing a respective lock bit for each of the plurality of constraint parameters on a reset event of the processor, wherein each lock bit disables the interface from changing a respective constraint parameter when set.

18. The non-transitory computer-readable storage medium of claim 17, wherein the interface of the processor permits the user and the application software to set the power management constraint during runtime of the processor without the reboot of the processor.

19. The non-transitory computer-readable storage medium of claim 17, wherein the power management constraint comprises a first set of constraints for when the processor is running a first application and a second set of constraints for when the processor is running a second application specified by the user.

20. The non-transitory computer-readable storage medium of claim 17, wherein the power management constraint comprises a first consumption target for when the processor is running a first application and a second consumption target for when the processor is running a second application specified by the user.

21. The non-transitory computer-readable storage medium of claim 17, wherein the power management constraint com-



prises a first set of constraints for when the processor is running a first combination of applications and a second set of constraints for when the processor is running a second combination of applications specified by the user.

22. The non-transitory computer-readable storage medium of claim 17, wherein the power management constraint comprises a first consumption target for when the processor is running a first combination of applications and a second consumption target for when the processor is running a second combination of applications specified by the user.

23. The non-transitory computer-readable storage medium of claim 17, the method further comprising clearing a lock bit on a reset event of the processor, wherein the lock bit disables the interface when set.

24. The non-transitory computer-readable storage medium of claim 1, wherein the plurality of constraint parameters are overclocking constraint parameters.

\* \* \* \* \*